

B1

LISTING OF CLAIMS

- 1.** (Currently Amended) A method for configuring a distributed storage system for a failover condition, said method comprising the steps of:

providing a plurality of distributed object storage managers "DOSMs" for receiving requests for files;

providing at least three intelligent storage nodes accessible to said DOSMs over a network **interconnecting the DOSMs and the storage nodes;**

storing at least one file in a first intelligent storage node accessed via a DOSM over said network;

storing a duplicate of said file in a second intelligent storage node accessed via said network;

entering a failover condition to cease use of said first intelligent storage node;

determining, at said DOSM, a location for said file in said second intelligent storage node;

identifying said second storage node as said location for said file; and
accessing, via said network, said file stored in said second intelligent storage node in response to a subsequent file request.

- 2.** (Original) The method as set forth in claim 1, wherein:

the step of storing at least one file in a first intelligent storage node accessed via a network comprises the step of accessing said first intelligent storage node via a first network address;

B1

the step of storing a duplicate of said file in a second intelligent storage node accessed via said network comprises the step of accessing said second intelligent storage node via a second network address; and

the step of determining a location for said file in said second intelligent storage node comprises the step of generating a mapping between said first network address and said second network address.

3. (Original) The method as set forth in claim 2, wherein said first network address and said second network address comprise internet protocol (“IP”) network addresses and differ only in a subnet portion of said IP network addresses.

4. (Original) The method as set forth in claim 1, wherein:

the step of storing at least one file in a first intelligent storage node comprises the step of storing said file in said first intelligent storage node located in a first storage center; and

the step of storing a duplicate of said file in a second intelligent storage node comprises the step of storing said file in said second intelligent storage node located in a second storage center, geographically distant from said first storage center.

5. (Original) The method as set forth in claim 4, further comprising the steps of:

storing a plurality of files in a plurality of intelligent storage nodes in said first storage center; and

storing duplicates of said files in a plurality of intelligent storage nodes in said second storage center and said intelligent storage nodes in said first storage and said intelligent storage nodes in said second storage center.

B | 6. (Original) The method as set forth in claim 1, wherein:

the step of storing at least one file in a first intelligent storage node comprises the step of storing said file in said first intelligent storage node located in a first storage center; and

the step of storing a duplicate of said file in a second intelligent storage node comprises the step of storing said file in said second intelligent storage node located in said first storage center.

7. (Original) The method as set forth in claim 1, wherein:

the steps of storing at least one file in a first intelligent storage node comprises the steps of:

storing said file in a first storage center comprising a plurality of intelligent storage nodes;

storing said duplicate of said file in a second storage center, geographically distant from said first storage center; and

the step of determining a location for said file in said second intelligent storage node comprises the step of searching for said file in said second storage after entering said failover condition.

8. (Previously Presented) The method as set forth in claim 7, further comprising the step of searching for said file using a point-to-point protocol between distributed object storage manager (DOSM) and said second intelligent storage node.

9. (Original) The method as set forth in claim 7, further comprising the step of searching for said file using a multi-cast protocol.

B1

10. (Original) The method as set forth in claim 7, further comprising the step of searching for said file using a point-to-point protocol between a distributed object storage manager (DOSM) and an intelligent storage node.

11. (**Currently Amended**) A distributed storage system comprising:

a network;

at least three intelligent storage nodes;

a first intelligent storage node having multiple storage devices, accessed via said network, for storing at least one file;

a second intelligent storage node having multiple storage devices, accessed via said network, for storing a duplicate of said file; and

a plurality of distributed object storage managers (DOSMs) remote from said storage nodes, one of said DOSMs for accessing, via an interconnection over said network, said file stored in said first intelligent storage node in response to a file request, for entering a failover condition to cease access to said file in said first intelligent storage node, for determining a location for said file in one of said intelligent storage nodes, for identifying said second storage node as said location for said file, and for accessing, via said interconnection over said network, said file stored in said second intelligent storage node in response to a subsequent file request.

12. (Original) The distributed storage system as set forth in claim 11, wherein said DOSM further accessing said first intelligent storage node via a first network address and accessing said second intelligent storage node via a second network address, said DOSM further comprising a mapping between said first network address and said network address.

B | 13. (Original) The distributed storage system as set forth in claim 12, wherein said first network address and said second network address comprise internet protocol (“IP”) network addresses and differ only in a subnet portion of said IP network addresses.

14. (Original) The distributed storage system as set forth in claim 11, wherein:
a first storage center comprising said first intelligent storage node; and
a second storage center, geographically distant from said first storage center comprising said second intelligent storage node.

15. (Original) The distributed storage system as set forth in claim 14, wherein:
said first storage center comprises a plurality of files stored in a plurality of intelligent storage nodes; and

said second storage center comprises duplicates of said files stored in a plurality of intelligent storage nodes, so as to provide a one to one mapping between said intelligent storage nodes in said first storage center and said intelligent storage nodes in said second storage center.

16. (Original) The distributed storage system as set forth in claim 11, wherein said first and second intelligent storage nodes reside in a single storage center.

17. (Original) The distributed storage system as set forth in claim 11, further comprising:
a first storage center comprising a plurality of intelligent storage nodes for storing said file;

a second storage center, geographically distant from said first storage center, for storing said duplicate of said file; and

said DOSM for searching for said file in said second storage center after entering said failover condition.

B | 18. (Original) The distributed storage system as set forth in claim 17, wherein said DOSM further comprising processes for searching for said file in said first storage center if said file is not located in said second storage center.

19. (Original) The distributed storage system as set forth in claim 17, further comprising a multi-cast protocol for communicating among said DOSM and said intelligent storage nodes.

20. (Original) The distributed storage system as set forth in claim 17, further comprising a point-to-point protocol for communicating between said DOSM and one of said intelligent storage nodes.

21. (**Currently Amended**) A distributed virtual file system comprising:

a network;

a first directory, remote from a requesting client and from an associated intelligent storage node, accessed via said network, for storing file system information associated with said intelligent storage node having multiple storage devices;

a second directory, accessed via said network, for storing a duplicate of said file system information; and

at least one distributed directory manager (DDM) for accessing, via said network, said file system information stored in said first directory in response to a file system request, for entering a failover condition to cease access to said file system information in said first directory, for determining a location for said file system information in said second directory, and for accessing, via said network, said file system information stored in said second directory in response to a subsequent file request.

22. (Original) The distributed virtual file system as set forth in claim 21, wherein said DDM further for accessing said first directory via a first network address and accessing said second

B | directory via a second network address, said DDM further comprising a mapping between said first network address and said second network address.

23. (Original) The distributed virtual file system as set forth in claim 22, wherein said first network address and said second network address comprise internet protocol (“IP”) network addresses and differ only in a subnet portion of said IP network addresses.

24. (Original) The distributed virtual file system as set forth in claim 21, wherein:

- a first storage center comprising said first directory; and
- a second storage center, geographically distant from said first storage center comprising said second directory.

25. (Original) The distributed virtual files system as set forth in claim 24, wherein:

- said first storage center comprises file system information stored in a plurality of directories; and
- said second storage center comprises a duplicate of said files system information stored in a plurality of directories, so as to provide a one to one mapping between said directories in said first storage center and said directories in said second storage center.

26. (Original) The distributed virtual file system as set forth in claim 21, wherein said first and second directories reside in a single storage center.

27. (Original) The distributed virtual file system as set forth in claim 21, further comprising:

- a first storage center comprising a plurality of directories for storing said file system information;
- a second storage center, geographically distant from said first storage center, for storing a duplicate of said file system information; and

B1
said DDM for searching for said file system information in said second storage center after entering said failover condition.

28. (Original) The distributed virtual file system as set forth in claim 27, wherein said DDM further comprising processes for searching for said file system information in said first storage center if said file system information is not located in said second storage center.

29. (Original) The distributed virtual file system as set forth in claim 27, further comprising a multi-cast protocol for communicating among said DDM and said directories.

30. (Original) The distributed virtual file system as set forth in claim 27, further comprising a point-to-point protocol for communicating between said DDM and one of said directories.
